

Timeout Control in Distributed Systems Using Perturbation Analysis: Multiple Communication Links

Ali Kebarighotbi and Christos G. Cassandras *

Division of Systems Engineering
and Center for Information and Systems Engineering
Boston University
Brookline, MA 02446
alik@bu.edu, cgc@bu.edu

Abstract

Timeout control is a simple mechanism used when direct feedback is either impossible, unreliable, or too costly, as is often the case in distributed systems. Its effectiveness is determined by a timeout threshold parameter and our goal is to quantify the effect of this parameter on the system behavior. In this paper, we extend previous results to the case where there are N transmitting nodes making use of a common communication link bandwidth. After deriving the stochastic hybrid model for this problem, we apply Infinitesimal Perturbation Analysis to find the derivative estimates of aggregate average goodput of the system. We also derive the derivative estimate of the goodput of a transmitter with respect to its own timeout threshold which can be used for local and hence, distributed optimization.

1. Introduction

Timeout control is a simple mechanism used in many systems where direct feedback is either impossible, unreliable, or too costly. This is often the case in distributed systems, where remote components cannot be observed by a controller (or several distributed controllers) and information is provided over a network with unreliable links and delays. A timeout event is scheduled using a *timer* which expires after some *timeout threshold* parameter. This defines an expected time by which some other event should occur or a “grace period” over which some response about the system state

is required. If no information arrives within this period, a “timeout event” occurs and incurs certain reactions which are an integral part of the controller. In distributed systems, where usually control decisions must be made with limited information from remote components, timeouts provide a key mechanism through which a controller can infer valuable information about the unobservable system states. In fact, as pointed out in [28], timeouts are indispensable tools in building up reliable distributed systems.

This simple reactive control policy based on timeouts has been used for stabilizing systems ranging from manufacturing to communication systems [14], Dynamic Power Management (DPM) [2],[13],[22] and software systems [7],[9] among others. Despite its wide usage, quantifications of its effect on system behavior have not yet received the attention they deserve. In fact, timeout controllers are usually designed based on heuristics which may lead to poor results; an important example can be found in communication protocols, especially TCP (see [1],[14],[21],[28] and references therein). There is limited work on finding optimal timeout thresholds. For instance, in [20], a single queueing model is used for an Automatic Guided Vehicle (AGV) system and shared tester equipment. In DPM where the aim is to minimize the average power consumption, [13] and [2] propose a timeout control scheme with the aid of the theory of competitive analysis. Also, in [3], a Markov process model is used and Infinitesimal Perturbation Analysis (IPA) [4],[10] is applied to calculate the optimal timeout threshold values. In communication systems, [11],[18],[19] have attempted to find the optimal TCP retransmission timeouts by making assumptions on the probability of the transmission failure in the system. Finally, in [12] optimal web session timeouts are calculated so as to reduce the probability of falsely ending a web session in time sensitive web pages. All

*The authors' work is supported in part by NSF under Grant EFRI-0735974, by AFOSR under grant FA9550-09-1-0095 and FA9550-12-1-0113, by DOE under grant DE-FG52-06NA27490, by ONR under grant N00014-09-1-1051, and by ARO under grant W911NF-11-1-0227.

such approaches are limited by their reliance on the distributional information about the stochastic processes involved.

In this paper, we consider the timeout control in the context of Discrete Event Systems (DES), so that the controller output includes a response to either a timeout event or an event carrying information about network congestion. However, since stochastic DES models can be very complicated to analyze, we rely on recent advances which abstract a DES into a Stochastic Hybrid System (SHS) and, in particular, the class of *Stochastic Flow Models* (SFM). A SFM treats the event rates as stochastic processes of arbitrary generality except for mild technical conditions. The emphasis in using SFMs is not in deriving approximations of performance measures of the underlying DES, but rather studying sample paths from which one can derive structural properties which are robust with respect to the abstraction made. This is the case, for instance, with many performance gradient estimates which can be obtained through IPA techniques for general SHS [6],[25]. In addition, a fundamental property of IPA in SFMs (as in DES) is that the derivative estimates obtained are independent of the probability laws of the stochastic rate processes and require minimal information from the observed sample path. This approach has proved useful in optimizing various performance metrics in serial networks [23], systems with feedback control mechanisms [27], scheduling problems [15],[16], and some multi-class models [5],[24].

In [17], we set forth a line of research aimed at quantifying how timeout threshold parameters affect the system state and ultimately its behavior and performance. We adopted a communication system model consisting of one transmitter submitting packets to a network with stochastic processing times. We showed how a SFM of a timeout-controlled distributed system can be obtained and then used to optimize the timeout threshold. This paper extends the results in [17] to multiple communication links and hence, extends this problem to broader multi-class problems with communication delays. Like [17], we directly control the timeout parameters for network communication performance. Additionally, moving from a single class problem in [17] to multiple classes, we show how by carefully selecting the processing rate of each class, the shared communication channel operates according to a First Come First Serve (FCFS) policy. We also discuss how we can easily extend the analysis to non-FCFS frameworks. Finally, we also extend previous work on timeout control by removing any dependence on distributional information.

This paper is organized as follows: In Section 2, we

define the SFM of the system. In Section 3, we apply IPA techniques to average goodput at each transmitter, as well as total goodput of the system aggregating the rates for all the transmitting nodes. We conclude with Section 4.

2. Stochastic Flow Model (SFM)

We are interested in systems where N transmitting nodes send data to their destinations through a shared channel. Associated with each node $n = 1 \dots, N$, is a “class” n of data packets or tasks which are to be processed in the channel and reach their destinations in a timely manner. Like the case of a single communication link, after each transmission by node n , a timely acknowledgement (ACK) is expected from the receiving end. For transmitter node n to operate normally, the *Round-Trip Transportation* (RTT) time calculated for each node n should be less than a *timeout threshold* parameter $\theta_n \geq 0$. If no ACK is received within θ_n units of time from the transmission, a *timeout* occurs causing node n to retransmit the timed out data (previously sent at $t - \theta_n$) to the same destination. The control parameter vector is defined as $\theta = [\theta_1, \theta_2, \dots, \theta_N]^T$. We con-

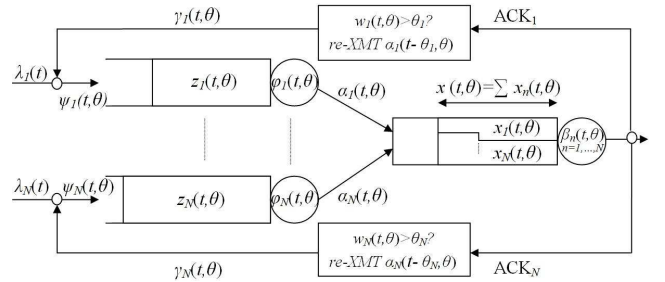


Figure 1. The SFM for the timeout system

sider a SFM for this problem where N upstream nodes are competing to get their content through a shared network channel. The channel latency represents the RTT delay in the underlying DES. The SFM for this system is shown in Fig. 1 and is observed over a finite period of time $[0, T]$. Associated with this system are several nonnegative stochastic processes all defined on a common probability space (Ω, \mathcal{F}, P) . The transmitters are shown as upstream nodes receiving exogenous inflow processes $\{\lambda_n(t)\}$, $n = 1, \dots, N$ and copy/timeout flow processes $\{\gamma_n(t, \theta)\}$ intended for retransmission. The process $\{z_n(t, \theta)\}$ models the amount of pending data

to be transmitted by node n and evolves according to

$$\dot{z}_n(t, \theta) = \frac{dz_n(t, \theta)}{dt^+} = \begin{cases} 0 & \text{if } z_n(t, \theta) = 0, \psi_n(t, \theta) \leq \phi_n(t, \theta) \\ \psi_n(t, \theta) - \phi_n(t, \theta) & \text{otherwise.} \end{cases} \quad (1)$$

where

$$\psi_n(t, \theta) = \lambda_n(t) + \gamma_n(t, \theta) \quad (2)$$

and $\{\phi_n(t, \theta)\}$ is the *maximal transmission rate* process at node n . Thus, the *actual transmission rate* from node n to the network buffer is defined as

$$\alpha_n(t, \theta) = \begin{cases} \psi_n(t, \theta) & \text{if } z_n(t, \theta) = 0, \psi_n(t, \theta) \leq \phi_n(t, \theta) \\ \phi_n(t, \theta) & \text{otherwise} \end{cases} \quad (3)$$

In this paper, we assume that the transmitters have an *infinite supply* property as follows:

Assumption 1. $z_n(t, \theta) > 0$ for all $t \in [0, T]$.

By (3), this implies $\alpha_n(t, \theta) = \phi_n(t, \theta) > 0$ for all $t \in [0, T]$, so $\phi_n(t, \theta)$ can henceforth be replaced by $\alpha_n(t, \theta)$ for the rest of the discussion.

We model the shared network channel with a buffer where the transmitted fluid from node n accumulates and takes a share $x_n(t, \theta) \geq 0$ of the total buffer content $x(t, \theta) = \sum_n x_n(t, \theta)$. The processing rate of the fluid is governed by nonnegative service processes $\{\beta_n(t, \theta)\}$, $n = 1, \dots, N$.

Furthermore, for each n , we define the independent process $\{B_n(t)\}$ as the maximal processing rate of class n when no other class is being processed by the network resource. Since, $\beta_n(t, \theta)$ is the processing rate in the presence of other fluid classes competing for a share of processing power in the network, we naturally have $\beta_n(t, \theta) \leq B_n(t)$.

At any time $t \in [0, T]$, $x_n(t, \theta)$ in the network follows the dynamics

$$\dot{x}_n(t, \theta) = \frac{dx_n(t, \theta)}{dt^+} = \begin{cases} 0 & \text{if } x(t, \theta) = 0, \alpha(t, \theta) \leq \beta(t, \theta) \\ \alpha_n(t, \theta) - \beta_n(t, \theta) & \text{otherwise} \end{cases} \quad (4)$$

where $\alpha(t, \theta) = \sum_n \alpha_n(t, \theta)$ and $\beta(t, \theta) = \sum_n \beta_n(t, \theta)$. We also define the process

$$w_n(t, \theta) = \min_{w \geq 0} \left\{ \int_{t-w}^t \beta_n(\tau, \theta) d\tau = x_n(t-w, \theta) \right\}, \quad (5)$$

as the *waiting time* process of class $n \in \{1, \dots, N\}$. We assume that the value of the processes $\beta_n(t, \theta)$ and

$x_n(t-w, \theta)$ are known over $t \in [-w, 0]$ as the initial pieces of information required to calculate $w(t, \theta)$ at $t = 0$. Definition (5) has a very close relation to the “time-to-empty” process as defined in [26],[8] since one way to interpret (5) is as the earliest time it takes the server to process the content at $t-w$.

Using the waiting time process, we can define the timed-out (hence, worthless) portion of class- n fluid in the network buffer for which $w_n(t, \theta) > \theta_n$. This portion needs to be retransmitted by the upstream node n . Under normal conditions, the network generates an acknowledgement flow indicating a successful transmission previously made at $t - w_n(t, \theta)$. In case of a time-out at node n - i.e. the moment $w_n(t, \theta)$ exceeds θ_n - a copy of the fluid sent at time $t - \theta_n$ is generated and will be put back in node n for retransmission. This is shown as the feedback flow process $\gamma_n(t, \theta) \geq 0$ in Fig. 1. To summarize, for all $t \in [0, T]$, we have

$$\gamma_n(t, \theta) = \begin{cases} 0 & \text{if } w_n(t, \theta) \leq \theta_n \\ \alpha_n(t - \theta_n, \theta) & \text{otherwise} \end{cases} \quad (6)$$

The upstream node treats this copy flow with higher priority and tries to retransmit it as soon as possible. This is inline with many retransmission-based communication protocols. Looking at Fig.1, we should explicitly show the flow $\gamma_n(t, \theta)$ creating contents ahead of those created by $\lambda_n(t)$. However, we refrained from this to avoid complicating the illustration.

We look at *average goodput* as the communication performance metric defined as

$$J(\theta, T) = E[G(\theta, T)] \quad (7)$$

$$G(\theta, T) = \sum_n \int_0^T [\alpha_n(t, \theta) - 2\gamma_n(t, \theta)] dt \quad (8)$$

where, like the one-dimensional problem [17], we use $2\gamma_n(t, \theta)$ to not only penalize the retransmissions but also to account for the worthless timed out fluid in the network buffer which should be processed by the network resource anyway.

Before proceeding, we make the following assumption on the system which we also made in [17]:

Assumption 2. The network buffer is lossless.

This assumption is merely for simplifying the exposition and can be removed at the expense of having new state variables tracking the loss volume of each class.

2.1. FCFS Implementation

The definition of the waiting time (5) is contingent upon having a FCFS policy implemented in the network buffer. One of the differences between a single transmitter model in [17] is in how to ensure this policy is

preserved. We achieve this by carefully selecting the network processes.

In general, depending on the processes $\beta_n(t, \theta)$, if two fluid particles of class n and m , $n \neq m$ are transmitted at times τ_1 and τ_2 , $\tau_2 > \tau_1$, if $\beta_m(t, \theta)$ is sufficiently larger than $\beta_n(t, \theta)$, their order of leaving the network may be reversed. We can implement both FCFS and non-FCFS policies by carefully defining the processing rates $\beta_n(t, \theta)$, $n = 1, \dots, N$. We limit the details on the non-FCFS policies to Remark 1 where it is shown that the extension of the results to the non-FCFS policies is straightforward.

Recall that $\beta_n(t, \theta) \leq B_n(t)$ since some part of the resource is consumed by other classes. Moreover, just as in the DES where a slow service for one type of customer increases the waiting time of the other customers behind it, if $B_n(t)$ is small, it not only means slow processing for class n , but also a decreasing effect on other processing rates $\beta_m(t, \theta)$, $m \neq n$. The *availability rate* of fluid class n at the network server can be defined as

$$\tilde{\alpha}_n(t, \theta) = \alpha_n(t - w_n(t, \theta), \theta) \quad \forall t \in [0, T]. \quad (9)$$

Accordingly, at any time $t \in [0, T]$, the utilization of the service associated with class $n = 1, \dots, N$ is as follows:

$$\rho_n(t, \theta) = \frac{\tilde{\alpha}_n(t, \theta)}{B_n(t)}. \quad (10)$$

Hence, at each time t , $\rho_n(t, \theta)$ is a measure of how engaged the server is with the fluid of class n . The following theorem ensures all these characteristics are delivered by carefully defining $\beta_n(t, \theta)$, $n = 1, \dots, N$. Compared to [26],[8], it allows for different maximal service processes for each class.

Theorem 1. Let $w \in \mathbb{R}^+$. Assuming the initial conditions $w_n(0, \theta) \equiv w$ for all n , if

$$\beta_n(t, \theta) = \frac{\tilde{\alpha}_n(t, \theta)}{\sum_m \rho_m(t, \theta)} = \frac{\tilde{\alpha}_n(t, \theta)}{\sum_m \frac{\tilde{\alpha}_m(t, \theta)}{B_m(t)}}, \quad \forall t \in [0, T], \quad (11)$$

the following statements are true:

- (i) There exists a common waiting time process $w(t, \theta)$ such that $w_n(t, \theta) = w(t, \theta)$ for all n and all $t \in [0, T]$.
- (ii) If $0 \leq t_1 < t_2$, then $t_1 - w(t_1, \theta) < t_2 - w(t_2, \theta)$

Proof: When $x(t, \theta) = 0$, the buffer clearly operates according to FCFS, so we only consider the case where $x(t, \theta) > 0$. By proving statement (i), we show the fluid particles at the head of the queue must have arrived at the same time; Statement (ii) shows that regardless of the class, fluid particles leave the system with the order

they have been transmitted.

We start by proving statement (i). Differentiating the term in brackets in (5) with respect to t reveals

$$\beta_n(t, \theta) - [1 - \dot{w}_n(t, \theta)]\beta_n(t - w_n(t, \theta), \theta) = [1 - \dot{w}_n(t, \theta)]\dot{x}_n(t - w_n(t, \theta), \theta)$$

which after regrouping the terms yields

$$\dot{w}_n(t, \theta) = 1 - \frac{\beta_n(t, \theta)}{\alpha_n(t - w_n(t, \theta), \theta)} = 1 - \frac{\beta_n(t, \theta)}{\tilde{\alpha}_n(t, \theta)}. \quad (12)$$

Using (11) in the last equation gives

$$\dot{w}_n(t, \theta) = 1 - \frac{1}{\sum_m \rho_m(t, \theta)}, \quad \forall t \in [0, T]. \quad (13)$$

for which we consider the constraint $w_n(t) \geq 0$. Notice that the RHS is independent of n . Hence, with the initial condition that $w_n(0, \theta) \equiv w$ for all n , $w_n(t, \theta) = w_m(t, \theta)$ for any pair n, m , $n \neq m$ and the FCFS operation when $w_n(t, \theta)$ is defined for all n . Finally, we can define the waiting time dynamics of a fluid differential at the head of the network buffer as

$$\dot{w}(t, \theta) = \begin{cases} 0 & \text{if } x(t, \theta) = 0 \\ 1 - \frac{1}{\sum_m \rho_m(t, \theta)} & \text{otherwise} \end{cases} \quad (14)$$

This proves part (i) in the theorem statement.

For part (ii) notice that by definition, a fluid particle which is at the head of the queue at time t has arrived at $t - w(t, \theta)$. Therefore, we only need to show that for any $t_1, t_2 \in [0, T]$ such that $t_1 < t_2$, we have $t_1 - w_1(t_1, \theta) < t_2 - w_2(t_2, \theta)$. We can write $t_1 - w(t_1, \theta) < t_2 - w(t_2, \theta)$. It follows that

$$\begin{aligned} t_1 - t_2 &< w(t_1, \theta) - w(t_2, \theta) \\ &= w(t_1, \theta) - \left[w(t_1, \theta) + \int_{t_1}^{t_2} \dot{w}(\tau, \theta) d\tau \right] \\ &= - \int_{t_1}^{t_2} \dot{w}(\tau, \theta) d\tau \\ &= t_1 - t_2 - \int_{t_1}^{t_2} \frac{1}{\sum_m \rho_m(\tau, \theta)} d\tau \end{aligned}$$

which, by (14), gives

$$0 < \int_{t_1}^{t_2} \frac{1}{\sum_m \rho_m(\tau, \theta)} d\tau.$$

Notice that $\frac{1}{\sum_m \rho_m(\tau, \theta)} > 0$. This proves part (ii) and the theorem. ■ With a more complicated proof, we can show the theorem statements

are true even if Assumption 1 does not apply. However, we do not consider it in this paper.

Remark 1: We can come up with schemes other than FCFS and generalize the SFM analysis. For example, if by defining

$$\beta_n(t, \theta) = B_n(t) \frac{\tilde{\alpha}_n(t, \theta)}{\sum_m \tilde{\alpha}_m(t, \theta)}, \quad \forall n.$$

we share the resource capacity according to relative availability of the fluid classes at the server, (12) becomes

$$\dot{w}(t) = 1 - \frac{B_n(t)}{\sum_m \tilde{\alpha}_m(t, \theta)}, \quad \forall t \in [0, T]. \quad (15)$$

Now, the right-hand side of this equation is n -dependent, which breaks the FCFS rule.

For the rest of this paper, we make the following assumption:

Assumption 3. $B_n(t) = B(t)$ for all $t \in [0, T]$ and $n = 1, \dots, N$.

We can remove this assumption at the expense of more complexity which diverts the focus from the main purpose of this analysis. Assumption 3 is not limiting in the present problem as the communication channels treat the packets from different sources or classes equally. Moreover, a byproduct of Assumption 3 is that

$$\begin{aligned} \beta(t, \theta) &= \sum_n \beta_n(t, \theta) = \sum_n \frac{\tilde{\alpha}_n(t, \theta)}{\sum_m \frac{\tilde{\alpha}_m(t, \theta)}{B_m(t)}} \\ &= B(t) \sum_n \frac{\tilde{\alpha}_n(t, \theta)}{\sum_m \tilde{\alpha}_m(t, \theta)} = B(t), \quad \forall t \in [0, T], \end{aligned} \quad (16)$$

which means the total processing rate of the channel is independent of the timeout rates chosen.

2.2. Transmission Control

We say node n is operating in a *normal period* when $w(t, \theta) \leq \theta_n$. In this mode, we assume that the inflow rates are determined according to the policy $\pi_{1,n}$ which is designed to increase the transmission rate $\alpha_n(t, \theta)$. We adopt a second policy $\pi_{2,n}$ which applies when the node is in the *timeout period*, i.e., $w(t, \theta) > \theta_n$. Policy $\pi_{2,n}$ is reactive and aims at reducing the transmission rate until the network channel comes out of the congestion and can satisfies the requirement $w(t, \theta) \leq \theta_n$ again. There can be many choices for the policies $\pi_{1,n}$ and $\pi_{2,n}$. For the sake of analysis, we choose the policies as follows:

$$\alpha_n(t, \theta) = f_{\alpha_n}(t, \theta) = \begin{cases} r_a & \text{if } w(t, \theta) \leq \theta_n \\ 0 & \text{otherwise} \end{cases}, \quad (17)$$

$$\alpha_n(t, \theta) \equiv \alpha_{n,\min} \text{ if } w(t, \theta) > \theta_n.$$

2.3. Stochastic Hybrid Model

Viewed as a SHS, we can conceive of the following SFM operation modes throughout the sample path: We refer to a period over which $x(t, \theta) > 0$ and $x(t, \theta) = 0$, as a *Non-Empty Period* (NEP) and *Empty Period* (EP), respectively. Moreover, we denote the periods over which $w(t, \theta) > \theta_n$ and $w(t, \theta) \leq \theta_n$ by TOP_n and NP_n , respectively.

Let τ_k , $k = 1, \dots, K$, be the SFM event times observed in a sample path of the system over the interval $[0, T]$. We also define $\tau_0 = 0$ and $\tau_{K+1} = T$ for notational convenience and let e_k be the event occurring at τ_k . We are interested in the following set of events:

$$\begin{aligned} \mathcal{E} = \{ & E_\lambda, E_B, [x > 0], [x = 0], [w > \theta_n], [w \leq \theta_n], \\ & [\tilde{\alpha}_n^+ \neq \tilde{\alpha}_n^-], [\gamma_n^+ \neq \gamma_n^-] \}. \end{aligned} \quad (18)$$

Here, E_λ and E_B respectively refer to random jump events in any $\lambda_n(t)$, $n = 1, \dots, N$ and $B(t)$ which are not affected by θ . We call the events with this property *exogenous*. The start and end of a NEP is marked by the events $[x > 0]$ and $[x = 0]$, respectively. $[w > \theta_n]$ and $[w \leq \theta_n]$ are the events defining the start and end of a TOP_n . Since these events are dependent on the system states, we categorize them as *endogenous* events. When a timeout $[w > \theta_n]$ event occurs, according to (17), $\alpha_n(t, \theta)$ drops. This discontinuity will then be reflected in $\tilde{\alpha}_n(t, \theta)$ and consequently, by (11), in $\beta_n(t, \theta)$. We refer to the event of $\tilde{\alpha}_n(t^+, \theta) \neq \tilde{\alpha}_n(t^-, \theta)$ by $[\tilde{\alpha}_n^+ \neq \tilde{\alpha}_n^-]$. As the last system event, $[\gamma_n^+ \neq \gamma_n^-]$ is an event of discontinuity in $\gamma_n(t, \theta)$ strictly inside a TOP . Since the start and stop of a TOP_n is already defined by $[w > \theta_n]$ and $[w \leq \theta_n]$, by (6), $[\gamma_n^+ \neq \gamma_n^-]$ only reflects a discontinuity in $\alpha_n(t - \theta_n, \theta)$ when $w(t, \theta) > \theta_n$. This directly affects the objective function (8) as it is related to $\gamma_n(t, \theta)$ defined by (6). We call $[\tilde{\alpha}_n^+ \neq \tilde{\alpha}_n^-]$ and $[\gamma_n^+ \neq \gamma_n^-]$ *induced* events because these events are bound to occur after the occurrence of a *triggering* event in the past. Generally, the triggering event can be either an exogenous, endogenous or itself an induced event. More information on the induced events can be found in [6].

The delays between $[w > \theta_n]$ and its associated $[\tilde{\alpha}_n^+ \neq \tilde{\alpha}_n^-]$ and $[\gamma_n^+ \neq \gamma_n^-]$ call for new state variables. The role of these state variables is to provide *timers* which trigger when $[w > \theta_n]$ occurs and measure the amount of time until the associated induced event. For

each event time τ_k we define the state variables

$$h_1(k, t, \theta) = \begin{cases} -\beta(t) & \text{if } \exists n : e_k = [w > \theta_n] \\ 0 & \text{otherwise} \end{cases} \quad \forall t \geq \tau_k, \quad (19a)$$

$$h_1(k, \tau_k^+, \theta) = \begin{cases} x(t_k, \theta) & \text{if } e_k = [w > \theta_n] \\ 0 & \text{otherwise} \end{cases} \quad (19b)$$

$$h_2(k, t, \theta) = \begin{cases} -1 & \text{if } \exists n : e_k = [w > \theta_n] \\ 0 & \text{otherwise} \end{cases} \quad \forall t \geq \tau_k, \quad (20a)$$

$$h_2(k, \tau_k^+, \theta) = \begin{cases} \theta_n & \text{if } e_k = [w > \theta_n] \\ 0 & \text{otherwise} \end{cases} \quad (20b)$$

with the constraint $h_i(k, t, \theta) \geq 0, i = 1, 2, k = 1, \dots, K$. To identify the active timers on $[\tau_k, \tau_{k+1})$, for each $k = 1, \dots, K$, we define the index set

$$\Phi_k = \{m \leq k : \exists i \text{ s.t. } h_i(m, t, \theta) > 0 \forall t \in [\tau_k, \tau_{k+1})\}. \quad (21)$$

3. Performance Optimization by IPA

The following assumption ensures existence of the IPA derivatives:

Assumption 4. With probability 1, no two events can occur at the same time unless one causes the other.

Considering the objective (8), let us define the index set

$$\Omega_n = \{k : w(t, \theta) > \theta_n \forall t \in [\tau_k, \tau_{k+1})\} \quad (22)$$

marking all the event times in TOP_n including its start. The other possible event in Ω_n is $[\tilde{\alpha}_n^+ \neq \tilde{\alpha}_n^-]$ which by Assumption 4, cannot occur independently of $[w > \theta_n]$ and $[w \leq \theta_n]$. Since by (6), when $w(t, \theta) > \theta_n$, we have $\gamma_n(t, \theta) = \alpha_n(t - \theta_n, \theta)$, we can decompose (8) as follows:

$$G(T, \theta) = \sum_n G_n(T, \theta), \quad (23a)$$

$$G_n(T, \theta) = \left\{ \sum_{k=0}^K \int_{\tau_k}^{\tau_{k+1}} \alpha_n(t, \theta) dt - 2 \sum_{k \in \Omega_n} \int_{\tau_k}^{\tau_{k+1}} \alpha_n(t - \theta_n, \theta) dt \right\}, \quad n = 1, \dots, N. \quad (23b)$$

3.1. IPA Estimation

Let us define $\tau'_{k,j} \equiv \frac{\partial \tau_k}{\partial \theta_j}$ for $k = 0, \dots, K+1$ and $j = 1, \dots, N$. Also, for a real valued function $f_n(t, \theta)$

associated with node n , we define the partial derivatives $f'_{n,j}(t, \theta) \equiv \frac{\partial f_n(t, \theta)}{\partial \theta_j}$. Differentiating $G_n(T, \theta)$ with respect to θ_j and noticing by $\tau_0 = 0$, and $\tau_{K+1} = T$, that $\tau'_{0,j} = \tau'_{K+1,j} = 0$ for any j reveals

$$\begin{aligned} \frac{dG_n(T, \theta)}{d\theta_j} &= \left\{ \sum_{k=1}^K [\alpha_n(\tau_k^-) - \alpha_n(\tau_k^+)] \tau'_{k,j} \right. \\ &\quad - 2 \sum_{k \in \Omega_n} \left[\tau'_{k+1,j} \alpha_n([\tau_{k+1} - \theta_n]^-) - \tau'_{k,j} \alpha_n([\tau_k - \theta_n]^+) \right. \\ &\quad \left. \left. + \sum_{k=0}^K \int_{\tau_k}^{\tau_{k+1}} \alpha'_{n,j}(t) dt + \int_{\tau_k}^{\tau_{k+1}} \frac{d\alpha_n(t - \theta_n)}{d\theta_j} dt \right] \right\}. \end{aligned} \quad (24)$$

The expression in the last integral in (24) can be written as

$$\frac{d\alpha_n(t - \theta_n, \theta)}{d\theta_j} = \frac{\partial \alpha_n(\tau, \theta)}{\partial \theta_j} \Big|_{\tau=t-\theta_n} + \frac{\partial(t - \theta_n)}{\partial \theta_j} \dot{\alpha}_n(t - \theta_n, \theta)$$

where $\frac{\partial \alpha_n(\tau, \theta)}{\partial \theta_j} \Big|_{\tau=t-\theta_n} \equiv \alpha'_{n,j}(\tau - \theta_n, \theta)$. Moreover, according to (17), $\dot{\alpha}_n(t - \theta_n, \theta) = r_n$ if $w_n(t - \theta_n, \theta) < \theta_n$ (i.e., node n not in a TOP_n at $t - \theta_n$). Noting that $\frac{\partial(t - \theta_n)}{\partial \theta_j} = -1$ only if $j = n$ and 0, otherwise, we find

$$\begin{aligned} \frac{d\alpha_n(t - \theta_n, \theta)}{d\theta_j} &= \alpha'_{n,j}(t - \theta_n, \theta) \\ &\quad - \begin{cases} r_n & \text{if } w(t - \theta_n, \theta) \leq \theta_n, j = n \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (25)$$

Looking at the performance objective as well as the conditions under which an event is triggered, the state vector of the system is comprised of $\alpha_n(t, \theta)$, $\gamma_n(t, \theta)$, $n = 1, \dots, N$, $x(t, \theta)$ and $w(t, \theta)$. However, note that $\gamma_n(t, \theta) = \alpha_n(t - \theta_n, \theta)$ when $w_n(t, \theta) > \theta_n$. Thus, its derivative can be obtained from that of $\alpha_n(t, \theta)$. Furthermore, Assumption 3, allows us to have the following useful lemma which reduces the number of states to only two:

Lemma 1. If the network channel operates according to the FCFS policy, we get

$$\frac{\partial w(t, \theta)}{\partial \theta_j} = \frac{1}{\tilde{\alpha}(t, \theta)} \frac{\partial \tilde{x}(t, \theta)}{\partial \theta_j}, \quad \forall t \in [0, T], j = 1, \dots, N,$$

where $\tilde{x}(t, \theta) = x(t - w(t, \theta), \theta)$ for all $t \in [0, T]$.

Proof: Differentiating (5) with respect to θ_j $j \in \{1, \dots, N\}$ gives

$$\begin{aligned} &w'_{n,j}(t, \theta) \beta_n(t - w_n(t, \theta), \theta) + \int_{t-w_n(t, \theta)}^t \frac{d\beta_n(\tau, \theta)}{d\theta_j} d\tau \\ &= \frac{\partial x_n(\tau, \theta)}{\partial \theta_j} \Big|_{\tau=t-w_n(t, \theta)} + \frac{\partial x_n(\tau, \theta)}{\partial \tau} \frac{\partial \tau}{\partial \theta_j} \Big|_{\tau=t-w_n(t, \theta)} \\ &= x'_{n,j}(t - w_n(t, \theta), \theta) - \dot{x}_n(t - w_n(t, \theta), \theta) w'_{n,j}(t, \theta) \end{aligned}$$

where we have $\dot{x}_n(t - w_n(t, \theta), \theta) = \alpha_n(t - w_n(t, \theta), \theta) - \beta_n(t - w_n(t, \theta), \theta)$. Thus,

$$\begin{aligned} w'_{n,j}(t, \theta) \beta_n(t - w_n(t, \theta), \theta) + \int_{t-w_n(t, \theta)}^t \frac{d\beta_n(\tau, \theta)}{d\theta_j} d\tau \\ = x'_{n,j}(t - w_n(t, \theta), \theta) \\ - [\alpha_n(t - w_n(t, \theta), \theta) - \beta_n(t - w_n(t, \theta), \theta)] w'_{n,j}(t, \theta) \end{aligned}$$

Regrouping terms yields the following for all $t \in [0, T]$:

$$w'_{n,j}(t, \theta) = \frac{x'_{n,j}(t - w_n(t, \theta), \theta) - \int_{t-w_n(t, \theta)}^t \frac{d\beta_n(\tau, \theta)}{d\theta_j} d\tau}{\tilde{\alpha}_n(t, \theta)}.$$

Using Theorem 1 and by Assumption 1, $w_n(t, \theta) = w(t, \theta)$ for all $n = 1, \dots, N$ and all $t \in [0, T]$. Since $w'_n(t, \theta) = w'(t, \theta)$ for all $n \in \{1, \dots, N\}$ we can add all the numerators and denominators and write

$$\frac{\partial w(t, \theta)}{\partial \theta_j} = \frac{1}{\sum_n \tilde{\alpha}_n(t, \theta)} \sum_n \left[x'_{n,j}(t - w(t, \theta), \theta) - \int_{t-w(t, \theta)}^t \frac{d\beta_n(\tau, \theta)}{d\theta_j} d\tau \right].$$

Since $\sum_n \tilde{\alpha}_n(t, \theta) = \tilde{\alpha}(t, \theta)$ and $\sum_n \frac{d\beta_n(t, \theta)}{d\theta_j} = \frac{d\beta(t)}{d\theta_j} = 0$, we find that for $j = 1, \dots, N$,

$$\frac{\partial w(t, \theta)}{\partial \theta_j} = \frac{1}{\tilde{\alpha}(t, \theta)} \frac{\partial x(t - w(t, \theta), \theta)}{\partial \theta_j} = \frac{1}{\tilde{\alpha}(t, \theta)} \frac{\partial \tilde{x}(t, \theta)}{\partial \theta_j}.$$

■

3.2. IPA equations

Before proceeding, we provide a brief review of the IPA framework for general stochastic hybrid systems as presented in [6]. If $s(t, \theta) \in \mathbb{R}^M$ is the state vector of the SFM, IPA specifies how changes in θ influence $s(t, \theta)$ and the event times τ_k and, ultimately, how they influence interesting performance metrics which are generally expressed in terms of these variables. Let us assume that over an interval $[\tau_k, \tau_{k+1})$, the SFM is at some mode during which the time-driven state satisfies $\dot{s} = f_k(s, \theta, t)$ for some $f_k : \mathbb{R}^M \times \mathbb{R}^N \times [0, T) \rightarrow \mathbb{R}^M$. Let $s'(t) \equiv \frac{\partial s(t)}{\partial \theta} \in \mathbb{R}^M \times \mathbb{R}^N$ be the Jacobian matrix for all state derivatives. It is shown in [6] that, for any $t \in [\tau_k, \tau_{k+1})$, $s'(t)$ satisfies:

$$\frac{d}{dt} s'(t) = \frac{\partial f_k(t)}{\partial s} s'(t) + \frac{\partial f_k(t)}{\partial \theta} \quad (26)$$

for $t \in [\tau_k, \tau_{k+1})$ with boundary condition:

$$s'(\tau_k^+) = s'(\tau_k^-) + \tau'_k [f_{k-1}(\tau_k^-) - f_k(\tau_k^+)] \quad (27)$$

for $k = 0, \dots, K$ if $s(t)$ is continuous at τ_k and otherwise,

$$s'(\tau_k^+) = r'(\tau_k). \quad (28)$$

An exogenous event at τ_k is not a function of θ so $\tau'_{k,j} = 0$ for any $j = 1, \dots, N$. In our model these include E_λ and E_B . However, for every endogenous event, e_k at τ_k there exists a continuously differentiable function $g_k : \mathbb{R}^n \times \Theta \rightarrow \mathbb{R}$ such that $\tau_k = \min\{t > t_{k-1} : g_k(s(t; \theta), \theta) = 0\}$. It is shown in [6] that

$$\tau'_k = - \left[\frac{\partial g_k}{\partial s} f_k(\tau_k^-) \right]^{-1} \left(\frac{\partial g_k}{\partial \theta} + \frac{\partial g_k}{\partial s} s'(\tau_k^-) \right) \quad (29)$$

if $e_k \in \mathcal{E}$ is endogenous and defined as long as $\frac{\partial g_k}{\partial s} f_k(\tau_k^-) \neq 0$. In addition to the exogenous and endogenous events, we also have induced events. Using the results in [6], in this case, if e_k is induced by e_m , $m < k$, we can write

$$\begin{aligned} \frac{\partial \tau_k}{\partial \theta_j} = - [h_i(m, \tau_k^-, \theta)]^{-1} \left(\frac{\partial s(\tau_k^-)}{\partial \theta_j} \frac{\partial h_i}{\partial s} \right. \\ \left. + \frac{\partial h_i(m, \tau_k^-, \theta)}{\partial \theta_j} \right). \quad (30) \end{aligned}$$

where $\frac{\partial h_i}{\partial s}$ is the gradient vector of partial derivatives h_i with respect to state variables and depending on whether $[\tilde{\alpha}_n^+ \neq \tilde{\alpha}_n^-]$ or $[\gamma_n^+ \neq \gamma_n^-]$ has occurred, we use h_1 or h_2 as defined in (19) and (20).

3.2.1. Event-time derivatives. We will determine $\frac{\partial \tau_k}{\partial \theta_j}$ for each event in \mathcal{E} and for each θ_j , $j = 1, \dots, N$. We exclude E_λ and E_B as they are exogenous events with zero event-time derivative. Recalling $\alpha(t) = \sum_n \alpha_n(t)$, $\tilde{\alpha}(t) = \sum_n \tilde{\alpha}_n(t)$ and $\beta(t) = \sum_n \beta_n(t)$, the following lemma gives the derivatives for exogenous and induced events:

Lemma 2. Under policies $\pi_{1,n}$ and $\pi_{2,n}$, $n = 1, \dots, N$, for any $j = 1, \dots, N$, we have

- (i) If $e_k = [x > 0]$: $\tau'_{k,j} = - \frac{1}{\alpha(\tau_k^-) - \beta(\tau_k^-)} \frac{\partial \alpha(\tau_k^-, \theta)}{\partial \theta_j}$
- (ii) If $e_k = [x = 0]$: $\tau'_{k,j} = - \frac{1}{\alpha(\tau_k^-, \theta) - \beta(\tau_k^-)} \frac{\partial x(\tau_k^-, \theta)}{\partial \theta_j}$
- (iii) If $e_k = [w > \theta_n]$ or $[w \leq \theta_n]$:
 $\tau'_{k,j} = \frac{\mathbf{1}_{n=j} \tilde{\alpha}(\tau_k, \theta) - \frac{\partial \tilde{x}(\tau_k^-, \theta)}{\partial \theta_j}}{\tilde{\alpha}(\tau_k^-, \theta) - \beta(\tau_k^-)},$
- (iv) If $e_k = [\tilde{\alpha}_n^+ \neq \tilde{\alpha}_n^-]$:
 $\tau'_{k,j} = \frac{1}{\beta(\tau_k^-)} \left[\frac{\partial x(\tau_m^+, \theta)}{\partial \theta_j} + \tau'_{m,j} \alpha(\tau_m, \theta) \right]$
- (v) If $e_k = [\gamma_n^+ \neq \gamma_n^-]$: $\tau'_{k,j} = \tau'_{m,j} + \mathbf{1}_{n=j}$

where $\mathbf{1}_{n=j}$ is an indicator function being 1 when $n = j$ and 0, otherwise and $\tau_m < \tau_k$ is the time of the triggering event for the induced event at τ_k .

Proof: Starting with part (i), we can invoke (29) with $g_k(t, \theta) = \alpha(t, \theta) - \beta(t)$. Noticing $\frac{\partial g_k}{\partial \alpha} = 1$ and $\frac{\partial g_k}{\partial \beta} = -1$, we find that $\frac{\partial g_k}{\partial \alpha} \dot{\alpha}(\tau_k^-, \theta) + \frac{\partial g_k}{\partial \beta} \dot{\beta}(\tau_k^-) = \dot{\alpha}(\tau_k^-, \theta) - \dot{\beta}(\tau_k^-)$. Moreover, $\frac{\partial g_k(\tau_k^-, \theta)}{\partial \theta_j} = \frac{\partial \alpha(\tau_k^-, \theta)}{\partial \theta_j} - \frac{\partial \beta(\tau_k^-)}{\partial \theta_j}$. However, by (16) we know that $\beta(t) = B(t)$ and is independent of the control parameters. Thus, we get $\frac{\partial g_k(\tau_k^-, \theta)}{\partial \theta_j} = \frac{\partial \alpha(\tau_k^-, \theta)}{\partial \theta_j}$. Inserting the results in the expression for (29) verifies part (i). For part (ii), notice that we have an endogenous event with $g_k(t, \theta) = x(t, \theta)$. Noticing that when $x(t, \theta) > 0$, $\dot{x}(t, \theta) = f_x(t, \theta) = \alpha(t, \theta) - \beta(t)$. Hence, $\frac{\partial g_k}{\partial x} f_x(\tau_k^-, \theta) = \alpha(\tau_k^-, \theta) - \beta(\tau_k^-)$. Also, $\frac{\partial g_k(\tau_k^-, \theta)}{\partial \theta_j} = \frac{\partial x(\tau_k^-, \theta)}{\partial \theta_j}$. These give $\frac{\partial \tau_k}{\partial \theta_j} = -\frac{1}{\alpha(\tau_k^-, \theta) - \beta(\tau_k^-)} \frac{\partial x(\tau_k^-, \theta)}{\partial \theta_j}$ which is exactly the claim of part (ii). For part (iii), we have $g_k(\tau_k, \theta) = w(\tau_k, \theta) - \theta_n$. By (13), and Assumption 3, we have

$$\dot{w}(t, \theta) = f_w(t, \theta) = 1 - \frac{1}{\sum_m \rho_m(t, \theta)} = 1 - \frac{B(t)}{\tilde{\alpha}(t, \theta)} = \frac{\tilde{\alpha}(t, \theta) - \beta(t)}{\tilde{\alpha}(t, \theta)}.$$

Since $\frac{\partial g_k}{\partial w} = 1$, we find $\frac{\partial g_k}{\partial w} f_w(\tau_k^-, \theta) = \frac{\tilde{\alpha}(\tau_k^-, \theta) - \beta(\tau_k^-)}{\tilde{\alpha}(\tau_k^-, \theta)}$. Finally, $\frac{\partial g_k(\tau_k^-, \theta)}{\partial \theta_j} = \frac{\partial w(\tau_k^-, \theta)}{\partial \theta_j} - \frac{\partial \theta_n}{\partial \theta_j}$. Clearly, $\frac{\partial \theta_n}{\partial \theta_j} = \mathbf{1}_{n=j}$. By Lemma 1, we get

$$\frac{\partial w(\tau_k^-, \theta)}{\partial \theta_j} = \frac{1}{\tilde{\alpha}(\tau_k^-, \theta)} \frac{\partial x([\tau_k - w(\tau_k, \theta)]^-, \theta)}{\partial \theta_j}.$$

Thus, we find that

$$\frac{\partial g_k(\tau_k^-, \theta)}{\partial \theta_j} = \frac{1}{\tilde{\alpha}(\tau_k^-, \theta)} \left(\frac{\partial x([\tau_k - w(\tau_k, \theta)]^-, \theta)}{\partial \theta_j} - \mathbf{1}_{n=j} \right).$$

Putting these results in (29) gives

$$\begin{aligned} \frac{\partial \tau_k}{\partial \theta_j} &= -\frac{1}{\frac{\partial g_k}{\partial w} f_w(\tau_k^-, \theta)} \frac{\partial g_k(\tau_k^-, \theta)}{\partial \theta_j} \\ &= \frac{\mathbf{1}_{n=j} \tilde{\alpha}(\tau_k, \theta) - \frac{\partial \tilde{\alpha}(\tau_k^-, \theta)}{\partial \theta_j}}{\tilde{\alpha}(\tau_k^-, \theta) - \beta(\tau_k^-)} \end{aligned}$$

which completes the proof for part (iii). For case (iv), we have the condition $h_1(\tau_k, \theta) = 0$ which assuming the triggering time is τ_m can be written as

$$h_1(m, \tau_k, \theta) = x(\tau_m, \theta) - \int_{\tau_m}^{\tau_k} \beta(t) dt = 0.$$

Notice that, $\dot{h}_1(\tau_k^-, \theta) = -\beta(\tau_k^-)$ and

$$\begin{aligned} \frac{\partial h_1(m, \tau_k^-, \theta)}{\partial \theta_j} &= \frac{\partial x(\tau_m, \theta)}{\partial \theta_j} + \frac{\partial \tau_m}{\partial \theta_j} \beta(\tau_m^+) \\ &= \frac{\partial x(\tau_m^-, \theta)}{\partial \theta_j} + \frac{\partial \tau_m}{\partial \theta_j} f_x(\tau_m^-, \theta) + \frac{\partial \tau_m}{\partial \theta_j} \beta(\tau_m^+). \end{aligned}$$

Since, at τ_m we have the event of discontinuity of α_n , and since by Assumption 4, E_B cannot coincide with this event, we have $\beta(\tau_m^-) = \beta(\tau_m^+)$. Using this in the above expression by knowing that $f_x(\tau_m^-, \theta) = \alpha(\tau_m^-, \theta) - \beta(\tau_m^-)$ reveals

$$\frac{\partial h_1(m, \tau_k^-, \theta)}{\partial \theta_j} = \frac{\partial x(\tau_m^-, \theta)}{\partial \theta_j} + \tau_{m,j}' \alpha(\tau_m^-, \theta).$$

Finally, notice that $h_1(m, \tau_k, \theta)$ is not directly dependent on the value of any state variable at τ_k , hence the first term in the paranthesized expression in (30) is 0. Putting all the results into (30) proves part (iv). Finally, for case (v), by (20), we have $h_2(m, t, \theta) = x(\tau_m, \theta) - (t - \tau_m)$ for $t \in [\tau_m, \tau_k)$ with the boundary condition $h_2(m, \tau_k, \theta) = 0$. Applying the same procedure proves part (v) and the whole lemma. ■

3.2.2. State Derivatives. Here, we find the derivative equations for the state variables x and α . We do not include other states in the IPA calculations as by (25) and Lemma 1, they can be calculated in terms of the derivatives of x and α :

A) Analysis at event times: We start by finding the derivative update equation of $\alpha_n(t, \theta)$. Notice that unlike [17] where $\alpha_n(t, \theta)$ was a continuous function of time, here, by definition $\alpha_n(t, \theta)$ has a discontinuity whenever event $[w_n > \theta_n]$ occurs. In general, we have the following lemma:

Lemma 3. Concerning state variable $\alpha_n(t, \theta)$,

$$\frac{\partial \alpha_n(\tau_k^+, \theta)}{\partial \theta_j} = 0, \text{ if } e_k = [w > \theta_n]. \quad (31)$$

Also, for the rest of the events, we have

$$\begin{aligned} \frac{\partial \alpha_n(\tau_k^+, \theta)}{\partial \theta_j} &= \frac{\partial \alpha_n(\tau_k^+, \theta)}{\partial \theta_j} \\ &+ \begin{cases} -r_n \frac{\mathbf{1}_{n=j} \tilde{\alpha}(\tau_k, \theta) - \frac{\partial \tilde{\alpha}(\tau_k^-, \theta)}{\partial \theta_j}}{\tilde{\alpha}(\tau_k^-, \theta) - \beta(\tau_k^-)} & \text{if } e_k = [w \leq \theta_n] \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (32)$$

Proof: Equation (31) is immediate from (28) noting $r(t) = \alpha_{n,\min}$ is independent of θ . Focusing on the second part, the proof of this lemma is also straightforward by using (27). Notice that the only event e_k (excluding $[w > \theta_n]$) at whose occurrence time the difference $f_{\alpha_n}(\tau_k^-, \theta) - f_{\alpha_n}(\tau_k^+, \theta)$ is non-zero is $e_k = [w \leq$

$\theta_n]$ where by (17), we get $f_{\alpha_n}(\tau_k^-, \theta) - f_{\alpha_n}(\tau_k^+, \theta) = 0 - r_n = -r_n$. Using this information in (27) proves the lemma. ■

Next, for any real-valued function f , we define $\Delta\alpha_n(\tau_k) = \alpha_n(\tau_k^-, \theta) - \alpha_n(\tau_k^+, \theta)$ for $k = 1, \dots, K$. We derive the discrete update equations for state variable x .

Lemma 4. Concerning state variable $x(t, \theta)$, we have

$$\begin{aligned} \frac{\partial x(\tau_k^+, \theta)}{\partial \theta_j} &= \frac{\partial x(\tau_k^-, \theta)}{\partial \theta_j} \\ &+ \begin{cases} -\frac{\partial x(\tau_k^-, \theta)}{\partial \theta_j} & \text{if } e_k = [x = 0] \\ \Delta\alpha_n(\tau_k, \theta) \frac{\mathbf{1}_{n=j} \tilde{\alpha}(\tau_k, \theta) - \frac{\partial \tilde{x}(\tau_k^-, \theta)}{\partial \theta_j}}{\tilde{\alpha}(\tau_k^-, \theta) - \beta(\tau_k^-)} & \text{if } e_k = [w > \theta_n] \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (33)$$

Proof: The proof follows by straightforward application of (27) to state variable x . When $e_k = [x = 0]$, we get $f_x(\tau_k^-, \theta) = \alpha(\tau_k^-, \theta) - \beta(\tau_k^-)$ and $f_x(\tau_k^+, \theta) = 0$. Hence, by using $\frac{\partial \tau_k}{\partial \theta_j}$ from Lemma 2, we find that

$$\begin{aligned} [f_x(\tau_k^-, \theta) - f_x(\tau_k^+, \theta)] \frac{\partial \tau_k}{\partial \theta_j} &= \\ [\alpha(\tau_k^-, \theta) - \beta(\tau_k^-)] \frac{-\frac{\partial x(\tau_k^-, \theta)}{\partial \theta_j}}{\alpha(\tau_k^-, \theta) - \beta(\tau_k^-)} &= -\frac{\partial x(\tau_k^-, \theta)}{\partial \theta_j}, \end{aligned}$$

which proves the first condition. When $e_k = [w > \theta_n]$, we again invoke Lemma 2 for this event and use it in (27) considering $f_x(\tau_k^-, \theta) - f_x(\tau_k^+, \theta) = \alpha_n(\tau_k^-, \theta) - \alpha(\tau_k^+, \theta)$. For other events, case by case analysis shows that $f_x(\tau_k^-, \theta) - f_x(\tau_k^+, \theta)$ is only non-zero when E_B occurs. However, this is an exogenous event with $\frac{\partial \tau_k}{\partial \theta_j} = 0$ for all $j = 1, \dots, N$. ■

B) Analysis between event times: We start from $\alpha_n(t, \theta)$. Recall that by (17) and Assumption 1, $f_{\alpha_n}(t, \theta) = f_{\phi_n}(t, \theta)$ for all n and $t \in [0, T)$. Since r_n is not a function of θ , differentiating $f_{\alpha_n}(t, \theta)$ with respect to θ_j reveals

$$\frac{\partial f_{\alpha_n}(t, \theta)}{\partial \theta_j} = 0, \quad \forall t \in [0, T) \text{ and } j = 1, \dots, N. \quad (34)$$

For $x(t, \theta)$, (4) implies that for all $t \in [0, T)$,

$$f_x(t, \theta) = \begin{cases} 0 & \text{if } x(t, \theta) = 0, \alpha(t, \theta) \leq \beta(t) \\ \alpha(t, \theta) - \beta(t) & \text{otherwise} \end{cases}$$

therefore, for $j = 1, \dots, N$,

$$\frac{\partial f_x(t, \theta)}{\partial \theta_j} = \begin{cases} 0 & \text{if } x(t, \theta) = 0, \alpha(t, \theta) \leq \beta(t) \\ \frac{\partial \alpha(t, \theta)}{\partial \theta_j} & \text{otherwise} \end{cases} \quad (35)$$

Using (25) and Lemma 1 and applying the IPA formulas given above completes the state derivative estimations.

3.2.3. IPA Implementation. We assume that the IPA derivatives for the states and event times over the interval $[t - w(t, \theta), t)$ are available for IPA calculations at t .

Recall that we are interested in estimating the derivative of the average cost (7) by finding the derivative of the sample function (8) which can be calculated from (24). However, evaluating (24) is contingent upon the knowledge of $\alpha'_{n,j}(t, \theta)$, $\frac{d\alpha_n(t - \theta_n, \theta)}{d\theta_j}$ and $\tau'_{k,j}$. By (25), $\frac{d\alpha_n(t - \theta_n, \theta)}{d\theta_j}$ can be readily obtained given that the past values of $\alpha'_{n,j}(t, \theta)$ and $w(t - \theta_n, \theta)$ are known. By Lemma 2, $\tau'_{k,j}$ can be obtained if the derivatives $\frac{\partial \alpha(t, \theta)}{\partial \theta_j}$, $\frac{\partial x(t, \theta)}{\partial \theta_j}$ and $\frac{\partial \tilde{x}(t, \theta)}{\partial \theta_j}$ are available. The latter is the value of $\frac{\partial x(t, \theta)}{\partial \theta_j}$ at $t - w(t, \theta)$. The derivative $\frac{\partial \alpha(t, \theta)}{\partial \theta_j}$ itself can be found by evaluating (33) at event times $[x = 0]$ and $[w > \theta_n]$, $n = 1, \dots, N$ and integrating (35) between these events. Furthermore, $\frac{\partial \alpha(t, \theta)}{\partial \theta_j}$ can be calculated by evaluating (31) at the occurrences of $[w > \theta_n]$ and (32) at those of $[w \leq \theta_n]$. Notice that by (34), no integration of $\frac{\partial f_{\alpha_n}(t, \theta)}{\partial \theta_j}$ is needed.

4. Conclusions

We extended the results in [17] to multiple communication links and laid down the conditions under which the network buffer works according to the FCFS policy. We also showed how extensions to non-FCFS policies are possible. Under the FCFS policy, we derived the derivative estimates of average goodput at each transmitter and used it to find the sensitivity of the aggregate goodput with respect to timeout thresholds.

References

- [1] M. Allman, V. Paxson, and W. Stevens. TCP congestion control. *RFC2581*, April 1999.
- [2] L. Cai and Y. H. Lu. Joint power management of memory and disk. In *Proceedings of the Conference on Design, Automation and Test in Europe. IEEE Computer Society*, pages 86–91, 2005.
- [3] X. R. Cao and H. F. Chen. Perturbation realization, potentials, and sensitivity analysis of markov processes. *IEEE Transactions on Automatic Control*, 42(10):1382–1393, oct. 1997.
- [4] C. G. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems*. Springer-Verlag New York, Inc, Secaucus, NJ, USA, 2006.

- [5] C. G. Cassandras, G. Sun, C. G. Panayiotou, and Y. Wardi. Perturbation analysis and control of two-class stochastic fluid models for communication networks. *IEEE Transactions on Automatic Control*, 48(5):770–782, 2003.
- [6] C. G. Cassandras, Y. Wardi, C. G. Panayiotou, and C. Yao. Perturbation analysis and optimization of stochastic hybrid systems. *European Journal of Control*, 16(6):642–664, 2009.
- [7] L. D. Catledge and J. E. Pitkow. Characterizing browsing strategies in the world-wide web. *Computer Networks and ISDN Systems*, 27(6):1065 – 1073, 1995.
- [8] M. Chen, J.-Q. Hu, and M. C. Fu. Perturbation analysis of a dynamic priority call center. *IEEE Trans. on Automatic Control*, 55:1191 – 1196, 2010.
- [9] R. Cooley, B. Mobasher, and J. Srivastava. Data Preparation for Mining World Wide Web Browsing Patterns. *Knowledge and Information Systems*, 1:5–32, 1999.
- [10] P. Glasserman. *Gradient Estimation via Perturbation Analysis*. Kluwer Academic Publishers, 1991.
- [11] T. Glatard, J. Montagnat, and P. X. Optimizing jobs timeouts on clusters and production grids. In *Seventh IEEE International Symposium on Cluster Computing and the Grid*, pages 100 – 107, may 2007.
- [12] D. He and A. Göker. Detecting session boundaries from web user logs. In *22nd Annual Colloquium on Information Retrieval Research*, pages 57–66, 2000.
- [13] S. Irani, S. Shukla, and R. Gupta. Online strategies for dynamic power management in systems with multiple power-saving states. *ACM Transactions Embedded Computer Systems*, 2(3):325–346, 2003.
- [14] R. Jain. Divergence of timeout algorithms for packet retransmissions. In *5th Annual International Phoenix Conference on Computers and Communications*, pages 174–179, March 1986.
- [15] A. Kebarighotbi and C. G. Cassandras. Revisiting the optimality of $c\mu$ rule with stochastic flow models. *48th IEEE Conference on Decision and Control*, pages 2304–2309, 2009.
- [16] A. Kebarighotbi and C. G. Cassandras. Optimal scheduling of parallel queues using stochastic flow models. *Journal of Discrete-Event Dynamic Systems*, 21:547–576, 2011.
- [17] A. Kebarighotbi and C. G. Cassandras. Timeout control in distributed systems using perturbation analysis. In *Proceedings of 50th IEEE Conference on Decision and Control*, pages 5437–5442, 2011.
- [18] A. Kesselman and Y. Mansour. Optimizing TCP retransmission timeout. In P. Lorenz and P. Dini, editors, *Lecture Notes in Computer Science- ICN*, volume 3421, pages 133–140. Springer Berlin, 2005.
- [19] L. Libman and A. Orda. Optimal retrial and timeout strategies for accessing network resources. *IEEE/ACM Transactions on Networking*, 10(4):551 – 564, Aug. 2002.
- [20] J. Peralta, P. Anussornnitisarn, and S. Nof. Analysis of a time-out protocol and its applications in a single server environment. *International Journal of Computer Integrated Manufacturing*, 16(1):1–13, 2003.
- [21] I. Psaras and V. Tsoussidis. Why TCP timers (still) don’t work well. *Computer Networks*, 51(8):2033–2048, 2007.
- [22] P. Rong. Determining the optimal timeout values for a power-managed system based on the theory of markovian processes: offline and online algorithms. In *Design, Automation and Test in Europe*, pages 1128–1133, 2006.
- [23] G. Sun, C. G. Cassandras, and C. G. Panayiotou. Perturbation analysis and optimization of stochastic flow networks. *IEEE Transactions on Automatic Control*, AC-49(12):2113–2128, 2004.
- [24] G. Sun, C. G. Cassandras, and C. G. Panayiotou. Perturbation analysis of multiclass stochastic fluid models. *Journal of Discrete Event Dynamical Systems*, 14(3):267–307, 2004.
- [25] Y. Wardi, R. Adams, and B. Melamed. A unified approach to infinitesimal perturbation analysis in stochastic flow models: The single-stage case. *IEEE Transactions on Automatic Control*, 55(1):89 –103, Jan 2010.
- [26] C. Yao and C. Cassandras. A solution of the lot sizing problem as a stochastic resource contention game. In *Decision and Control, 49th IEEE Conference on*, pages 6728 –6733, dec. 2010.
- [27] H. Yu and C. G. Cassandras. Perturbation analysis and feedback control of communication networks using stochastic hybrid models. *Journal of Nonlinear Analysis*, 65(6):1251–1280, 2006.
- [28] L. Zhang. Why TCP timers don’t work well. In *Proceedings of the ACM SIGCOMM Conference on Communications Architecture and Protocol*, pages 397–405, Stowe, Vermont, USA, 1986.